

SPECIFICATION

Electronic Version 1.2.8

Stylesheet Version 1.0

Fibre-Channel Arbitrated-Loop Split Loop Operation

Cross Reference to Related Applications

This application relates to co-pending Irish application no. S2000/0711 filed on September 7, 2000 entitled "Cross-Point Switch for a Fibre Channel Arbitrated Loop" naming Aedan Diarmid Cailean Coffey as inventor; and to co-pending Irish application no. S2000/0706 filed on September 7, 2000 entitled "Data Gathering Device for a Rack Enclosure" naming Aedan Diarmid Cailean Coffey et al as inventors.

Background of Invention

- [0001] The present invention relates to an enclosure services processor card arranged to selectively split a fibre-channel arbitrated-loop (FC-AL).
- [0002] Referring now to Figure 1, a backplane 10 suitable for being housed in a 19" 3U rack/shelf is shown schematically. For clarity, detail such as tracking or peripheral devices such as memory chips have been omitted where this detail is not necessary for describing the present invention.
- [0003] The backplane further includes a pair of edge connectors (not shown), each adapted to receive one of a potential pair of plug-in enclosure services processor cards 4 (only one shown), each including a respective processor 7. The processor may, for example, operate according to the SCSI Enclosure Services (SES), SCSI Accessed Fault-Tolerant Enclosure (SAF-TE), SCSI Enclosure Services (SES) or Intelligent Platform Management Interface (IPMI) enclosure management protocols.
- [0004] One or more of a range of SES Processor cards 4 can be plugged in and populated to provide advanced features, such as Ethernet, Direct Attach Fibre or Simple Network

Management Protocol (SNMP). As described in related patent application entitled "Data Gathering Device for a Rack Enclosure" naming Aedan Diarmid Cailean Coffey et al as inventors (Attorney Docket No. PI29273), the two processors 7 arbitrate between one another as to which will be responsible for the provision of enclosure services.

[0005] When one processor takes responsibility, by default, the other processor goes into a "dead till enabled" standby mode. If the responsible processor for some reason fails, then the other processor will take over responsibility for environmental monitoring.

[0006] A further plurality of edge connectors (not shown) are mounted on the backplane, and a set of Fibre Channel disk drives 80 are inserted into associated slots in the rack system where they connect at their rear to respective edge connectors and are thus interconnected via the backplane. Such disks include two pairs of fibre channel ports, each pair connecting to a respective one of a pair of Fibre Channel-Arbitrated Loops (FC-AL) designated Loop A and Loop B. Each disk further includes an Enclosure Services Interface (ESI) port (also more properly known as SFF-8067. SFF = Small Form Factor). This is conventionally used to provide communication between one or more of the disks and an enclosure services processor located on either a card connected to the backplane or even incorporated on the backplane.

[0007] A host application running on one or both of host machines Host A or Host B communicates with each of the disks 80 via one of the fibre channel loops. When the host application wishes to send and receive enclosure services information, for example, using the SES protocol, it sends a request to a pre-determined one of the disks. When the receiving disk processes the request, it determines that the request is an enclosure services request and forwards this request through its ESI port. Alternatively, the SES processor can report status to and receives commands from the host application out-of-band via an RS-232 connection (not shown).

[0008] In any case, for each disk 80, each of the sending and receiving ports for loop A and loop B are brought onto a respective processor card 4 where they connect to a hub 6. Conventionally, the hub comprises a plurality of port bypass circuits (PBC) each employed to connect a respective disk in a loop and to bypass ports where no disk is inserted in the slot. These circuits enable the fibre channel arbitrated loop (FC-AL) to

be opened and closed and thereby dynamically reconfigured if a device is added or removed to the FC-AL. The PBC comprises a 2-input multiplexer and a switch connecting either one of the inputs to the output. The PBCs are connected in sequence such that the output of each PBC acts as one of the inputs for the next PBC. The loop is completed by connecting the output of the last device in the loop, in this case the Host A to one of the inputs of the first PBC (PBC₁).

[0009] For example, looking at hub port PBC₂, it can be seen that the output of the preceding hub port, namely PBC₁, is transmitted both to hub port PBC₂ and to the input of Disk 2. The response from Disk 2 is transmitted back through the hub port PBC₂. In use, if a hub port (in this case PBC₂), detects that a device is connected to it, then the PBC directly connects the response signal from the attached device to the PBC output. For clarity, the detection circuitry is not shown for all circuits, however, in each case the output signal from a device whose presence is being tested is detected.

[0010] If however, a hub port detects that a device is not connected to it, or is not responding, it directly connects the input from the preceding PBC to its own output. This acts to close the associated PBC and bypass the hub port thereby allowing the loop to remain intact. This prevents a failing device or connection from bringing down the entire loop.

[0011] If it is desired to add a new device, at port PBC₁ for example, the loop opens automatically to add the new device without manual intervention, by toggling the position of the corresponding PBC switch. By this, the hub allows hot plugging; the ability to add and remove devices while the loop is active. In future versions of FC-AL, before a new device is allowed to be inserted in the loop, the hub will, at a minimum, verify valid signal quality. If a device exhibits poor signal quality or inappropriate clock speed, the associated PBC switch will remain toggled to bypass the hub-port, thereby allowing the other nodes on the loop to continue without disruption.

[0012] Arbitrated loop hubs may provide from 1 to 16 hub ports, i. e. n=16, with accommodation for more devices accomplished by cascading hubs together. A cascade is built by simply connecting a hub port, in this case PBC_{EX} of one hub to a hub port on the other, preferably with fibre optic cabling. Alternatively, a second host,

of port bypass circuits, each port bypass circuit being connected to a pair of tracks which in use connect to a respective one of each of said devices comprising said fibre channel arbitrated loop, said hub further comprising a pair of switches operatively controlled by said first switch, said pair of switches being disposed between respective port bypass circuits at which said loop is to be split so that in a first state said pair of switches connect said devices in a single loop and in a second state said pair of switches divide said devices into two split loops.

Brief Description of Drawings

- [0018] Embodiments of the invention will now be described, by way of example, with reference to the accompanying drawings, in which:
- [0019] Figure 1 is a schematic diagram of a enclosure services processor card including a conventional hub mechanism; and
- [0020] Figure 2 is a schematic diagram of an enclosure services processor card according to a preferred embodiment of the invention.

Detailed Description

- [0021] Referring now to Figure 2, in a preferred embodiment, Figure 2, a hub 6" of the type including port bypass circuits includes two additional switches or multiplexers MUX_{S1} and MUX_{S2} under the control of a switch 20 whose state is in turn controlled by the SES processor 7.
- [0022] In this case, it should be noted that because the loop is to be split, it is important that whenever one SES processor fails, the other processor can receive SES requests sent to at least one of the disks on its loop. This is a particularly important consideration when the loop may have been split into two loops prior to an SES processor card failure and may then need to be re-joined to allow access to the devices on a single loop via the hub on the other SES processor card.
- [0023] In the preferred embodiment, Disks 1 to n are connectable to the backplane and the SES processor 7 is configurable to selectively split its associated loop by actuating the multiplexers accordingly. This causes the loop to be split into two loops

comprising Disk 1 to Disk $n/2$; and Disk $n/2+1$ to Disk n respectively. Thus, MUX_{S1} is disposed between the output of $PBC_{n/2}$ and the input of $PBC_{n/2+1}$. The second input of MUX_{S1} is connected to the return signal from the expansion port to which in this case Host B is connected. MUX_{S2} is disposed between the output of PBC_{EX} and the input of Host A. The second input of MUX_{S2} is connected to the output signal from $PBC_{n/2}$. The remaining connections are as before.

[0024] Under normal non-split non-expanded loop operation the path around the loop comprises whichever are present of Disk 1 ... Disk $n/2$, through MUX_{S1} non-split input to whichever are present of Disk $n/2+1$... Disk n , through PBC_{EX} non-expanded input, through MUX_{S2} non-split input, through to Host A and back to Disk 1. It should be seen that when in this mode, as the output from $PBC_{n/2}$ is connected to the split input of MUX_{S2} , it can be regarded as open circuited.

[0025] The only difference between this and normal non-split expanded loop operation is that PBC_{EX} switches to its expanded input under the control of the detection circuitry 22, to allow in this case, Host B, to enter the loop between Disk n and Host A.

[0026] In split loop operation, the path around the first loop comprises whichever are present of Disk 1 ... Disk $n/2$, through MUX_{S2} split input through to Host A and back to Disk 1. It should be seen that when in this mode, as the output from $PBC_{n/2}$ is connected to the non-split input of MUX_{S1} , it can be regarded as open circuited.

[0027] The path around the second loop comprises whichever are present of Disk $n/2+1$... Disk n , probably through PBC_{EX} expanded input to include Host B and through MUX_{S1} split input to go back to Disk $n/2+1$. Again, it should be seen that when in this mode, as the output from PBC_{EX} is connected to the non-split input of MUX_{S2} , it can be regarded as open circuited.

[0028] It can therefore be seen that, if for example, the loop shown is Loop A and this is split, Host A can access Disk 1 to Disk $n/2$ on this loop and Host B can access Disk $n/2+1$ to Disk n on this same loop. However, maximum bandwidth is available to both hosts and so potentially double the transfer rate is available for the shelf.

[0029] If on Loop B, the hosts are added to the opposite split loop, i. e. Host A accessing

Disk $n/2+1$ to Disk n and Host B accessing Disk 1 to Disk $n/2$, then the throughput of the shelf can be almost doubled without any loss of access from the hosts to the disks.

[0030] Alternatively, up to four hosts can connect to respective split loops on each of Loop A and Loop B to potentially quadruple the throughput to disks vis-à-vis single host access.

[0031] Nonetheless, should a single host fail or should an enclosure services processor card fail, then the responsible SES processor can independently and automatically re-join the loops. Alternatively, any hosts that remain can instruct the SES processor via any disks it remains in contact with (or alternatively out-of-band) to re-join the loops.

[0032] It will be seen that the preferred embodiment has been described in terms of a SES processor 7 controlling the switch, which in turn controls the splitting of loop operation. Nonetheless, the invention could equally well be implemented by including a jumper on the SES processor card 4, with the jumper being open or closed to control the state of the switches joining or splitting the loop.

[0033] It will also be seen that with the addition of further switches such as MUX_{S1} and MUX_{S2} to the hub, the loop can be split into even more split loops as required.